


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide



THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

 Terms used **SAN backup archive partial restore**

Found 40 of 158,639

Sort results by

Display results

☒ [Save results to a Binder](#)
☒ [Search Tips](#)
☐ [Open results in a new window](#)
[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Results 1 - 20 of 40

 Result page: [1](#) [2](#) [3](#) [next](#)

 Relevance scale ☐ ☐ ☐ ☐ ☐

1 [ARIES: a transaction recovery method supporting fine-granularity locking and partial rollbacks using write-ahead logging](#)

C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz

 March 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 1

 Full text available: [pdf \(5.23 MB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

DB2TM, IMS, and TandemTM systems. ARIES is applicable not only to database management systems but also to persistent object-oriented languages, recoverable file systems and transaction-based operating systems. ARIES has been implemented, to varying degrees, in IBM's OS/2TM Extended Edition Database Manager, DB2, Workstation Data Save Facility/VM, Starburst and QuickSilver, and in the University of Wisconsin's EXODUS and Gamma d ...

Keywords: buffer management, latching, locking, space management, write-ahead logging

2 [Industrial sessions: beyond relational tables: Coordinating backup/recovery and data consistency between database and file systems](#)

Suparna Bhattacharya, C. Mohan, Karen W. Brannon, Inderpal Narang, Hui-I Hsiao, Mahadevan Subramanian

 June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

 Full text available: [pdf \(1.44 MB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


Managing a combined store consisting of database data and file data in a robust and consistent manner is a challenge for database systems and content management systems. In such a hybrid system, images, videos, engineering drawings, etc. are stored as files on a file server while meta-data referencing/indexing such files is created and stored in a relational database to take advantage of efficient search. In this paper we describe solutions for two potentially problematic aspects of such a data ...

Keywords: DB2, content management, database backup, database recovery, datalinks

3 [Peer-to-peer infrastructure: Pastiche: making backup cheap and easy](#)

Landon P. Cox, Christopher D. Murray, Brian D. Noble

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Full text available:  [pdf\(1.65 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#)

Backup is cumbersome and expensive. Individual users almost never back up their data, and backup is a significant cost in large organizations. This paper presents *Pastiche*, a simple and inexpensive backup system. Pastiche exploits excess disk capacity to perform peer-to-peer backup with no administrative costs. Each node minimizes storage overhead by selecting peers that share a significant amount of data. It is easy for common installations to find suitable peers, and peers with high ove ...

4 [DLFM: a transactional resource manager](#)

Hui-I Hsiao, Inderpal Narang

May 2000 **ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international conference on Management of data**, Volume 29 Issue 2

Full text available:  [pdf\(124.99 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The DataLinks technology developed at IBM Almaden Research Center and now available in DB2 UDB 5.2 introduces a new data type called DATALINK for a database to reference and manage files stored external to the database. An external file is put under a database control by "linking" the file to the database. Control to a file can also be removed by "unlinking" it. The technology provides transactional semantics with respect to linking or unlinking the file when DATALINK ...

5 [High speed on-line backup when using logical log operations](#)

David B. Lomet

May 2000 **ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international conference on Management of data**, Volume 29 Issue 2

Full text available:  [pdf\(220.69 KB\)](#)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Media recovery protects a database from failures of the stable medium by maintaining an extra copy of the database, called the backup, and a media recovery log. When a failure occurs, the database is "restored" from the backup, and the media recovery log is used to roll forward the database to the desired time, usually the current time. Backup must be both fast and "on-line", i.e. concurrent with on-going update activity. Conventional online backup sequentially copies ...

6 [The Recovery Manager of the System R Database Manager](#)

Jim Gray, Paul McJones, Mike Blasgen, Bruce Lindsay, Raymond Lorie, Tom Price, Franco Putzolu, Irving Traiger

June 1981 **ACM Computing Surveys (CSUR)**, Volume 13 Issue 2


Full text available:  [pdf\(1.75 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

7 [Principles of transaction-oriented database recovery](#)

Theo Haerder, Andreas Reuter

December 1983 **ACM Computing Surveys (CSUR)**, Volume 15 Issue 4

Full text available:  [pdf\(2.48 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

8 [Compiler transformations for high-performance computing](#)

David F. Bacon, Susan L. Graham, Oliver J. Sharp

December 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 4

Full text available:  pdf(6.32 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

In the last three decades a large number of compiler transformations for optimizing programs have been implemented. Most optimizations for uniprocessors reduce the number of instructions executed by the program using transformations based on the analysis of scalar quantities and data-flow techniques. In contrast, optimizations for high-performance superscalar, vector, and parallel processors maximize parallelism and memory locality with transformations that rely on tracking the properties of ...

Keywords: compilation, dependence analysis, locality, multiprocessors, optimization, parallelism, superscalar processors, vectorization

9 [File servers for network-based distributed systems](#)

Liba Svobodova

December 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 4Full text available:  pdf(4.23 MB)Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

10 [An efficient and flexible method for archiving a data base](#)


C. Mohan, Inderpal Narang

June 1993 **ACM SIGMOD Record , Proceedings of the 1993 ACM SIGMOD international conference on Management of data**, Volume 22 Issue 2Full text available:  pdf(969.13 KB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We describe an efficient method for supporting incremental and full archiving of data bases (e.g., individual files). Customers archive their data bases quite frequently to minimize the duration of data outage. Because of the growing sizes of data bases and the ever increasing need for high availability of data, the efficiency of the archive copy utility is very important. The method presented here minimizes interferences with concurrent transactions by not acquiring any locks on the data base ...

11 [Recovery Techniques for Database Systems](#)

Joost S. M. Verhofstad

June 1978 **ACM Computing Surveys (CSUR)**, Volume 10 Issue 2Full text available:  pdf(2.32 MB)Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

12 [A recovery algorithm for a high-performance memory-resident database system](#)

Tobin J. Lehman, Michael J. Carey

December 1987 **ACM SIGMOD Record , Proceedings of the 1987 ACM SIGMOD international conference on Management of data**, Volume 16 Issue 3Full text available:  pdf(1.50 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

With memory prices dropping and memory sizes increasing accordingly, a number of researchers are addressing the problem of designing high-performance database systems for managing memory-resident data. In this paper we address the recovery problem in the context of such a system. We argue that existing database recovery schemes fall short of meeting the requirements of such a system, and we present a new recovery mechanism which is designed to overcome their shortcomings. The proposed mechanism ...

13 [Operating systems: DualFS: a new journaling file system without meta-data duplication](#)

Juan Piernas, Toni Cortes, José M. García

June 2002 **Proceedings of the 16th international conference on Supercomputing**

Full text available:  [pdf\(213.64 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


In this paper we introduce DualFS, a new high performance journaling file system that puts data and meta-data on different devices (usually, two partitions on the same disk or on different disks), and manages them in very different ways. Unlike other journaling file systems, DualFS has only one copy of every meta-data block. This copy is in the *meta-data device*, a log which is used by DualFS both to read and to write meta-data blocks. By avoiding a time-expensive extra copy of meta-data b ...

Keywords: DualFS, journaling file system, meta-data management

14 Columns: Risks to the public in computers and related systems

Peter G. Neumann

January 2001 **ACM SIGSOFT Software Engineering Notes**, Volume 26 Issue 1

Full text available:  [pdf\(3.24 MB\)](#) Additional Information: [full citation](#)

15 Data collections and MM: 3D MURALE: multimedia database system architecture

Edward Grabczewski, John Cosmas, Peter Van Santen, Damian Green, Takebumi Itagaki, Fred Weimer

November 2001 **Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage**

Full text available:  [pdf\(169.47 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Archaeological databases are required to store a wide range of data about archaeological objects. Multimedia, spatial and temporal requirements are placing new demands on these databases. Virtual models of archaeological sites require new storage and search facilities, including searching of 3D graphics for virtual and physical restoration of archaeological finds. We examine the architecture, design philosophy and proposed implementation of the 3D MURALE multimedia database, which will be used b ...

Keywords: archaeology, multimedia Databases, photogrammetry, virtual Reality

16 Data base directions: the next steps

John L. Berg

November 1976 **ACM SIGMOD Record , ACM SIGMIS Database**, Volume 8 , 8 Issue 4 , 2

Full text available:  [pdf\(9.95 MB\)](#) Additional Information: [full citation](#), [abstract](#)

What information about data base technology does a manager need to make prudent decisions about using this new technology? To provide this information the National Bureau of Standards and the Association for Computing Machinery established a workshop of approximately 80 experts in five major subject areas. The five subject areas were auditing, evolving technology, government regulations, standards, and user experience. Each area prepared a report contained in these proceedings. The proceedings p ...

Keywords: DBMS, auditing, cost/benefit analysis, data base, data base management, government regulation, management objectives, privacy, security, standards, technology assessment, user experience

17

Industry perspectives: Toward autonomic computing with DB2 universal database

Sam S. Lightstone, Guy Lohman, Danny Zilio
September 2002 **ACM SIGMOD Record**, Volume 31 Issue 3


Full text available:  [pdf\(785.28 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

As the cost of both hardware and software falls due to technological advancements and economies of scale, the cost of ownership for database applications is increasingly dominated by the cost of people to manage them. Databases are growing rapidly in scale and complexity, while skilled database administrators (DBAs) are becoming rarer and more expensive. This paper describes the self-managing or autonomic technology in IBM's DB2 Universal Database® for UNIX and Windows to illustrate how self ...

18 ARIES/CSA: a method for database recovery in client-server architectures

C. Mohan, Inderpal Narang

May 1994 **ACM SIGMOD Record**, **Proceedings of the 1994 ACM SIGMOD international conference on Management of data**, Volume 23 Issue 2

Full text available:  [pdf\(1.33 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents an algorithm, called ARIES/CSA (Algorithm for Recovery and Isolation Exploiting Semantics for Client-Server Architectures), for performing recovery correctly in client-server (CS) architectures. In CS, the server manages the disk version of the database. The clients, after obtaining database pages from the server, cache them in their buffer pools. Clients perform their updates on the cached pages and produce log records. The log records are buffered locally ...

19 Relational Database Systems

Won Kim

September 1979 **ACM Computing Surveys (CSUR)**, Volume 11 Issue 3

Full text available:  [pdf\(2.67 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

20 Session 6: Log write-ahead protocols and IMS/VS logging

R. J. Peterson, J. P. Strickland

March 1983 **Proceedings of the 2nd ACM SIGACT-SIGMOD symposium on Principles of database systems**

Full text available:  [pdf\(2.14 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#)





Keywords: database, database management system, process synchronization point, recovery strategy, resource consistency, system failure, system log, transaction

Results 1 - 20 of 40

Result page: [1](#) [2](#) [3](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	8	(James near2 Ohr).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:26
L2	397	711/161.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:26
L3	1138	711/162.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:26
L4	37897	"709"/\$.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:27
L5	25755	"711"/\$.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:27
L6	1315	2 or 3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:30
L7	184503	(backup or (back\$3 adj up))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:31
L8	325985	(restor\$4 or (partial adj restor\$4) or unrestor\$4 or ("not" adj2 restor\$3))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:35
L9	1689921	map or table	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:35

L10	10575	7 and 8 and 9	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:35
L11	545	7 same 8 same 9	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:36
L12	181	7 with 8 with 9	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:36
L13	28	12 and 6	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:36
L14	73	11 and 6	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:37
L15	280827	SAN	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:37
L16	10	13 AND 15	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:37
L17	24	14 AND 15	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:37

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	8	(James near2 Ohr).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:26
L2	397	711/161.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:26
L3	1138	711/162.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:26
L4	37897	"709"/\$.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:27
L5	25755	"711"/\$.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:27
L6	1315	2 or 3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:30
L7	184503	(backup or (back\$3 adj up))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:31
L8	325985	(restor\$4 or (partial adj restor\$4) or unrestor\$4 or ("not" adj2 restor\$3))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:14
L9	1689921	map or table	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:35

L10	10575	7 and 8 and 9	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:35
L11	545	7 same 8 same 9	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:36
L12	181	7 with 8 with 9	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:36
L13	28	12 and 6	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:36
L14	73	11 and 6	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:37
L15	280827	SAN	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:37
L16	10	13 AND 15	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:37
L17	24	14 AND 15	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 13:37
L18	265	(unrestor\$4 or ("not" near3 restor\$3))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 15:01
L19	0	17 and 18	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:11

L20	0	16 and 18	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:11
L21	5	6 and 18	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:11
L22	70	(restor\$4 adj2 remainder\$2)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:43
L23	0	6 and 22	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:15
L24	0	("not" adj2 restor\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:17
L25	0	("not" adj5 restor\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:17
L26	0	("not" near2 restor\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:18
L27	0	('not' near2 restor\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:18
L28	430	(partial near2 restor\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:21
L29	10	6 and 28	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:21

L30	535	(restor\$4 adj2 remain\$5)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 15:08
L31	8	6 and 30	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:52
L32	0	("not" with restor\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:51
L33	69	"capability" adj restore	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:52
L34	265	(unrestor\$4)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:53
L35	5	6 and 34	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 14:53
L36	137	(incomplet\$4 or imcomplet\$4) near3 restor\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 15:02
L37	2	36 and 6	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/08/03 15:02